

DIGITAL SIGNAL PROCESSOR WITH PARALLEL ARCHITECTURE

Field of the Invention

The present invention relates to digital signal processors (DSPs), and more particularly, to a DSP with a parallel architecture or a VLIW (very large 5 instruction word) processor.

Background of the Invention

As is well known to those skilled in the art, VLIW type processors are processors derived from RISC (reduced instruction set computer) processors which 10 differ from conventional DSPs in that they comprise several parallel-mounted execution units. Each execution unit is the equivalent of a RISC processor core and executes instruction codes in reduced format, generally 16-bit codes, by exploiting the resources 15 offered by a bank of registers. Since each execution unit is capable of carrying out an instruction code simultaneously with the other execution units, the VLIW processors are therefore capable of simultaneously executing a large instruction comprising several RISC 20 equivalent codes in parallel.

To give a better understanding, Figure 1 presents a schematic view of the standard structure of a VLIW type processor 10 whose essential elements are shown in block form. The processor 10 comprises a 25 program memory PMEM, a data memory DMEM, an instruction register IR positioned at the output of the memory

PMEM, an instruction decoder IDEC positioned at the output of the register IR, a bank of registers REGBANK designed to execute the RISC type instruction codes, execution units EU0 to EU3, as well as a circuit BMC

5 forming the interface between the execution units UE1-EU3 and the inputs/outputs of the data memory DMEM. The execution units, which herein are four units EU0, EU1, EU2, EU3, are parallel-connected to simultaneously process four instruction codes that are read

10 simultaneously in the memory PMEM, together forming a large instruction. The nature of the execution units may vary as a function of the application for which the processor is designed. The execution units comprise for example an ALU (arithmetic and logic unit), a MAC

15 (multiplication/addition) unit, a CU (control unit managing the program counter PC and the connections), and a CO-PRO (coprocessor) unit to perform certain computations specific to the application.

A processor of this kind is thus capable of

20 executing large instructions which herein includes at most four codes. At each new clock cycle H, the program counter PC of the processor is increased by an increment n which is equal to 1, except in the case of a jump or a call, and the instruction registers IR0-IR3

25 receive four new codes simultaneously and in parallel. These four new codes are to be executed by the units EU0-EU3.

The architecture of a processor 10 of this kind thus differs from a conventional RISC processor by

30 its parallelism which can be found at all stages in the processing of the instructions. However, the possibilities offered by this parallelism are rarely exploited to the maximum, and the compiled programs stored in the program memory PMEM generally comprise a

35 large number of no-operation or NOP codes. Indeed, the conversion of a program written in a high-level language, for example the language C/C++, into a

sequence of RISC type codes combined in bundles is done automatically by a compilation program that knows the structure of the processor and tries to form bundles of the largest possible size (with a maximum of four codes
5 in the exemplary processor being described) to exploit the parallelism of the processor to the maximum. This optimization is done by taking account of the conflicts between the codes, the availability of the execution units and the data dependence during the pipeline
10 execution of the codes. Thus, for example, two codes designed to be executed by the same execution unit cannot be executed in parallel in the same bundle. Equally, a code using an operand that is the result of an operation that is object of another code cannot be
15 executed so long as the code on which it depends is not itself executed.

For example, let us consider the following program sequence:

/p c0	/p c1	p c2	/p c3	p c4	p c5	p c6	/p c7	p c8	p c9
-------	-------	------	-------	------	------	------	-------	------	------

20 which comprises instruction codes c0 to c9 comprising a parallelism bit /p or p. The instruction codes, known to those skilled in the art as syllables, are put together in bundles to form large instructions. The separation of the instructions (bundles) within a
25 program is done by the parallelism bits p assigned to each of the codes. The two possible values /p or p of a parallelism bit, for example 0 and 1, tells whether or not the code belongs to a new instruction. More particularly, a code preceded by a parallelism bit p
30 (for example 1) belongs to the same instruction as the previous code, while a code preceded by a parallelism bit /p (for example 0) belongs by convention to a new bundle.

In the program sequence mentioned above, the parallelism bits are thus used to distinguish four large instructions INST1 to INST4:

INST1 = c0
5 INST2 = c1 c2
INST3 = c3 c4 c5 c6
INST4 = c7 c8 c9

So that they can be executed by the processor 10, these instructions INST1 to INST4 are recorded in 10 the program memory PMEM of the processor as shown in Figure 1 and described in the following Table 1:

TABLE 1

/p c0	NOP	NOP	NOP
/p c1	NOP	NOP	NOP
15 /p c3	p c4	p c5	p c6
/p c7	p c8	p c9	NOP

Consequently, the compilers for standard VLIW processors generate a large number of no-operation codes for keeping certain execution units inactive 20 while others execute codes. This may result in the program memory becoming over-burdened with 20% to 70% of the memory space being taken up, depending on the efficiency of the compiler and the matching between the program to be compiled and the resources offered by the 25 processor. This over-burdening of the memory space by NOP codes, which is considerable in relation to the instruction programs really needed for the execution of the programs, increases the surface of the memory PMEM for a given application. This therefore causes a drop 30 in performance (the memory is slower with the increase in size), additional consumption (the word lines and bit lines are longer in the memory array) and costs (in terms of silicon surface area).

Summary of the Invention

In view of the foregoing background, an object of the present invention is to overcome the above described drawbacks. More particularly, the 5 object of the present invention is to provide a VLIW type processor capable of receiving, in its program memory, bundles of codes that do not comprise any no-operation codes.

These objects, features and advantages are 10 provided by a signal processor designed to execute variable-sized instructions that may comprise up to N elementary instruction codes. The processor comprises a memory program comprising I individually addressable, parallel-connected memory banks, with I being at least 15 equal to N, and wherein the codes of a program are recorded in an interlaced fashion at the rate of one code per bank and per address applied to the bank. The processor further comprises means for reading the program memory arranged to read a code in each of the I 20 memory banks during a cycle for reading an instruction. A cycle for reading an instruction in the program memory comprises reading a sequence of codes that comprises the instruction code or codes to be read and can also, when the number of codes of the instruction 25 read is smaller than I, comprise codes belonging to a following instruction.

According to one embodiment, the read means comprises means for the application, to the memory banks, of the individual addresses generated from a 30 collective value of a program counter that is incremented by a value equal to the number of codes comprising the previous instruction before the beginning of a cycle for reading an instruction.

According to another embodiment, the means to 35 apply addresses apply, to each of the memory banks, an individual read address equal to P0 or P0+1, with P0

being the quotient of the division by I of the value of the program counter.

According to another embodiment, the means for applying the addresses applies to an i ranking 5 memory bank an address equal to P0 when i is strictly greater than R, or for the applying to an i ranking memory bank an address equal to P0+1 when i less than or equal to R, with R being the remainder of the division by I of the value of the program counter.

10 According to yet another embodiment, the read means comprises means for reorganizing the codes of the sequence of codes read in the program memory, according to the following algorithm: $c'(j) = c(i)$, with $i = (j+R') \text{ modulo } I$, in which i and j designate the ranks of 15 the codes before and after reorganization, $c(i)$ designates i ranking codes in their arrangement after reading in the memory, $c'(j)$ designates j ranking codes after reorganization, and R' is the remainder of the division by I of the value that was shown by the 20 program counter during the previous clock cycle.

According to another embodiment, the means for reorganizing are arranged for the application to the sequence of codes read in the program memory of a circular permutation comprising a number of elementary 25 circular permutations equal to R' or to $I-R'$ depending on the direction in which the circular permutation is made.

According to yet another embodiment, the circular permutations are performed by a barrel shifter 30 receiving the parameter R' at a control input.

Advantageously, the read means comprises means to filter the codes that do not belong to the instruction to be read, using parallelism bits accompanying the codes.

35 According to another embodiment, the filtered codes are replaced by no-operation codes.

According to one embodiment, the code-filtering means are arranged to execute the following algorithm: For $j = 0$, $\text{val}(j=0) = v$, $s(j=0) = c'(j=0)$; For j going from 1 to I , $\text{val}(j) = v$ if $\text{val}(j-1) = v$ and
5 if parallelism bit of $c'(j) = p$, else $\text{val}(j-1) = /v$,
 $s(j) = c'(j)$ if $\text{val}(j) = v$; $s(j) = \text{NOP}$ if $\text{val}(j) = /v$,
in which $\text{val}(j)$ is a validation term associated with
each j ranking code, $c'(j)$ is capable of having two
values v and $/v$, $s(j)$ designates j ranking outputs of
10 the filtering means corresponding to same ranking
inputs receiving a code $c'(j)$, and NOP is a no-
operation code.

According to another embodiment, the non-
filtered codes are sent to parallel-mounted RISC type
15 execution units.

The present invention also relates to a
method for the reading of variable-sized instructions
that may include up to N elementary instruction codes,
applicable to a signal processor. The method comprises
20 the steps of providing for a program memory comprising
 I individually addressable parallel-connected memory
banks, with I being at least equal to N . The method
further includes recording the codes of a program in
the program memory in interlaced fashion, at a rate of
25 one code per bank and per address applied to the bank.
During a read cycle of an instruction, the method
includes reading a sequence of codes in the I memory
banks. The said sequence comprises the code or codes
of the instruction to be read and possibly also
30 comprises, when the number of instructions codes read
is smaller than I , with the codes belonging to a
following instruction.

According to yet another embodiment, the
method comprises the application, to the memory banks,
35 of the individual addresses generated from a collective
value of a program counter, that is incremented by a
value equal to the number of codes contained in the

previous instruction before the beginning of a cycle for reading an instruction.

According to another embodiment, the method comprises the application, to each of the memory banks, 5 of an individual read address equal to P0 or P0+1, with P0 being the quotient of the division by I of the value of the program counter. An i ranking memory bank receives an address equal to P0 when i is strictly greater than R or an address equal to P0+1 when i is 10 smaller than or equal to R, with R being the remainder of the division by I of the value of the program counter.

According to another embodiment, the method comprises the reorganization of the codes of the 15 sequence of codes read in the program memory, according to the algorithm mentioned above.

According to yet another embodiment, the method comprises a step of filtering the codes read, that do not belong to the instruction to be read, using 20 parallelism bits accompanying the codes.

According to another embodiment, filtered codes are replaced by no-operation codes.

According to another embodiment, the codes are filtered according to the algorithm mentioned 25 above.

According to yet a further embodiment, the non-filtered codes are sent to parallel-connected RISC type execution units.

Brief Description of the Drawings

30 These objects, features and advantages as well as others of the present invention shall be explained in greater detail in the following description of an exemplary embodiment of a processor according to the invention, given by way of a non- 35 restrictive description with reference to the appended figures, of which:

Figure 1, is a block diagram of the architecture of a standard VLIW processor, according to the prior art;

Figure 2 is a block diagram of the blocks of
5 the architecture of a processor according to the present invention; and

Figures 3 and 4 respectively illustrate detailed embodiments of selected elements represented in Figure 2.

10 **Detailed Description of the Preferred Embodiments**

Figure 2 is a schematic view of the architecture of a signal processor 20 according to the invention. The processor 20, like the standard processor of Figure 1, has a program memory PMEM1, a
15 data memory DMEM, an instruction register IR, an instruction decoder IDEC positioned at the output of the register IR, a register bank RBANK, execution units, in this case four units EU0 to EU3. The processor 20 further includes a BMC (built-in self-test/mapping/conflicts) circuit designed to read or write data elements in the data memory DMEM or in the register bank RBANK, and to fulfil various other functions such as the built-in self-test (BIST) function and the management of address conflicts. The
20 instruction register IR comprises four parallel-connected registers IR0 to IR3 designed for the reception, at each instruction reading cycle (or fetch cycle), of four instruction codes c'(0) to c'(3) delivered by the program memory PMEM1. The instruction
25 decoder IDEC decodes the codes c'(0)-c'(3) as well as distributes the operands to the execution units concerned. These elements are standard and form the core of the processor schematically shown in block form.
30

35 The processor 20 according to the invention differs from a standard processor in that the program

memory PMEM1 includes several parallel-connected memory banks, in this case four individually addressable banks M0, M1, M2, M3 that respectively receive addresses a0, a1, a2, a3. Hereinafter, the reference I designates the 5 number of individually addressable memory banks present in a processor according to the invention, N designates the maximum number of instruction codes that may be contained in a large instruction (bundle) and, I is at least equal to N. In the illustrated example, I is 10 equal to N which is equal to 4 since the processor comprises four execution units.

According to the invention, the instruction codes (syllables) of a program are recorded in a memory PMEM1 in an interlaced fashion, in such a way that two 15 codes succeeding each other in the program are arranged in two contiguous memory banks of increasing rank with the same address a_i in each bank. When a code is registered in the last memory bank M3, at an address $a_3 = x$, the next code is registered in the memory bank M0 20 at the address $a_0 = x+1$, and the following codes are registered in the banks M1, M2, M3 at the addresses $a_1 = a_2 = a_3 = x+1$, etc.

To provide a better illustration of the present invention, we shall assume that the following 25 program of sequences, already described in the introduction, is stored in the virtual memory PMEM1:

```
INST1 = c0  
INST2 = c1 c2  
INST3 = c3 c4 c5 c6  
30 INST4 = c7 c8 c9
```

Unlike a standard processor, the codes forming the large instructions INST1 to INST4 are recorded in the memory banks one after the other without any injection of no-operation codes NOP as can 35 be seen in the Table 2 below. The separation of the instructions is provided exclusively by the parallelism bits.

TABLE 2

Addresses	Bank M0	Bank M1	Bank M2	Bank M3	
5	X	/p c0	/p c1	p c2	/p c3
	x+1	p c4	p c5	p c6	/p c7
	x+2	p c8	p c9	/p...	...

It may be noted here that the parallelism bits are represented herein as elements external to the codes c0 to c9 to facilitate the understanding of the invention. In practice, the parallelism bits form an integral part of the codes stored in the memory and are eliminated at the time of the execution of the codes. For example, the codes c0-c9 comprise 16 bits, among which one bit is used as a parallelism bit.

Thus, advantageously, the processor according to the invention receives compact programs that do not comprise any NOP codes. The result thereof is a substantial reduction of the amount of memory space required. This results in a reduction of the surface area needed for the storage of the program and a reduction of consumption.

A description shall now be given for reading the memory PMEM1, designed to disinterlace the instruction codes in order to extract therefrom the large instructions or bundles, comprising several codes or syllables designed to be executed simultaneously.

According to the invention, these read means essentially comprise an address management circuit IAC, a code permutation circuit PCC, and a code selection circuit CSC and includes a special mode of management for the program counter PC of the processor 20 which shall be described first of all.

The program counter PC is incremented conventionally at each clock cycle, for example by an adder ADD1 receiving the current value of the program

counter at one input and an increment n at another input. The output of the adder is applied to a buffer LT1 driven by the clock signal H. According to the invention, the program counter PC is incremented by a value n which is not constant as in the prior art. The increment n is equal to the number of executable instruction codes contained in the previous instruction and is delivered by the circuit CSC described below.

This rule of incrementation of the program counter relates naturally to the sequential reading of the instructions of a program and cannot be applied during the execution of special instructions such as jump or call instructions for which the jump increment or the call value is dictated by the program.

The address management circuit IAC computes the individual addresses a₀-a₃ applied to the memory banks from the current value of the program counter PC, incremented in the manner that has just been described. More particularly, the circuit IAC computes the addresses a₀-a₃ by applying the following algorithm A1:

A1: for i ranging from 0 to I-1,
 a_i = P₀+1 if i < R
 a_i = P₀ if i = R

where:

1) i is the index designating the rank of the memory banks M₀-M₃ and the rank of the corresponding addresses a₀-a₃;
2) P₀ is the quotient of the division by I of the program counter PC, giving:

30 P₀ = Quotient (PC/I)

namely in this case I is equal to 4:

$$P_0 = \text{Quotient } (PC/4)$$

3) R is the remainder of the division by I of the program counter PC, giving:

$$R = PC - P_0 = PC \text{ modulo } I,$$

namely, in this case, I is equal to 4:

5 $R = PC \text{ modulo } 4$

Examples of the implementation of this algorithm as well as an exemplary embodiment of the circuit IAC shall be described further below. The code permutation circuit PCC is connected to the output of a
10 program memory PMEM1 and comprises inputs e0 to e3 respectively connected to the outputs of the memory banks M0 to M3 and respectively receives codes $c(0)$, $c(1)$, $c(2)$, $c(3)$ read in the memory program. The circuit PCC also comprises four outputs s0, s1, s2, s3
15 respectively delivering codes $c'(0)$, $c'(1)$, $c'(2)$, $c'(3)$. The codes delivered are equal to the codes received at input but are reorganized by the circuit PCC according to the following algorithm A2:

A2: for i ranging from 0 to -1
20 and j ranging from 0 to -1
 $c'(j) = c(i)$, with $i = (j + R') \text{ modulo } I$

in which:

- 1) i is the index, herein designating the rank of the input codes $c(0)-c(3)$ delivered by the memory banks M0-M3,
25
- 2) j is an index designating the rank of the output code $c'(0)-c'(3)$; and
- 3) R' is the remainder of the division by I of the program counter PC, computed at the previous clock
30 cycle.

As shown in Figure 2, the circuit PCC may be made by a conventional type barrel shifter. The algorithm A2 is implemented by the application to the control input of the barrel by a buffer LT2, of the 5 parameter R computed and delivered by the circuit IAC. Since the buffer LT2 is driven by the clock signal, the parameter R that is applied to the barrel shifter after each leading edge of the clock signal is shifted by a clock signal with reference to the parameter R of the 10 cycle in progress, and is therefore equal to the parameter R'. The barrel shifter PCC is positioned to shift its outputs in relation to its inputs by leftward circular permutation (in the direction going from the most significant values to the least significant 15 values), and by performing a number of leftward shifts equal to R'. In one alternative embodiment, the barrel shifter receives at its control input the complement of R' at I, namely the value [I-R'], and is laid out so as to shift its outputs in relation to its inputs by 20 rightward circular permutation (in the direction going from the least significant values to the most significant values) by performing a number of rightward shifts equal to [I-R'].

The code-selection circuit CSC mentioned 25 above is connected at its input to the output of the circuit PCC and comprises inputs e0 to e3 respectively receiving the codes $c'(0)$ to $c'(3)$. The circuit CSC also comprises four outputs s0 to s3 that copy the code received at the same ranking input in conditions set by 30 the following algorithm A3:

```
A3: for j = 0
    val(j=0) = v
    s(j=0) = c'(j=0)
    for j ranging from 1 to I
35        val(j) = v
        if val(j-1) = v
```

```
and if parallelism bit of c'(j)= p  
else val(j-1) = /v  
s(j) = c'(j) if val(j) = v,  
s(j) = NOP if val(j) = /v
```

- 5 where val(j) is a daisy chain type of recursive validation term associated with each j ranking output s(j) that can have two values v and /v, for example 1 and 0. For the j ranking outputs ranging from 1 to I, this term can be used to determine whether an output
10 s(j) must copy the code c'(j) received at the corresponding ranking input or, on the contrary, deliver a NOP code. This choice is a function of the value of the validation term val(j) which itself is a function of the value of the preceding validation term
15 val(j-1) and of the value p or /p, for example 1 or 0, of the parallelism bit present in the code c'(j).

In other words, assigning, by convention, the value 1 to p and v and the value 0 to /p and /v, the output s0 copies the code c'(0) received at the input
20 e0 and each of the other outputs s(j) copies the code c'(j) received at the corresponding input e(j) if the code c'(j) comprises a parallelism bit equal to 1, and if the validation bit val(j-1) of the previous ranking output is also equal to 1. When the parallelism bit of
25 a code c'(j) is equal to 0, the corresponding output s(j) and all the following outputs deliver a NOP code.

The system for reading the program memory PMEM1 thus works as follows:

Step E1: during a first clock cycle, the
30 circuit IAC applies addresses a0-a3 to the memory banks M0-M3, providing for the reading of four instruction codes a(0)-a(3) which follow each other in the memory PMEM1 according to the interlacing described above.

Step E2-1: during the following clock cycle,
35 a read signal READ is applied to the memory banks and the codes are read in an order which, because of their

interlacing, cannot correspond to their effective rank planned in the program. They are rearranged in order by the circuit PCC applying the algorithm A2 by the parameter R' computed in a previous clock cycle (where
5 the addresses a0-a3 were prepared).

Step E2-2: during the same clock cycles, the selection circuit CSC identifies the codes that do not correspond to the instruction to be executed, by means of their parallelism bit /p and replaces them with no-
10 operation codes NOP.

Step E2-3: during the same clock cycle, the instructions register IRO receives the code c'(0) and the registers IR1-IR3 receive the filtered codes (NOP) or non-filtered codes depending on the selection made
15 by the circuit CSC.

Step E3 and the following steps: the non-filtered codes are decoded by the decoder IDEC and executed by the units EU0-EU4.

Thus, the reading of a large instruction is
20 done in two clock cycles as in the prior art, and comprises a cycle for preparing the addresses a0-a3 and a cycle for the reading of the codes. The codes read are reorganized and filtered during the read cycle.
The operation is transparent for the operation of the
25 entire processor. These steps are performed with overlapping, according to the standard rules of pipeline operation. For the sake of simplicity, the pipeline registers have not been described or shown.

Table 3 shown in the appendix is an integral
30 part of the description and illustrates the operation of the processor for the reading of the instructions INST1 to INST4 described above. The parameters PC, R, R' and P0 that come into play in the execution of the algorithms are mentioned in the first columns. It is
35 considered in this example that the codes of the instructions INST1 to INST3 are stored in the program memory PMEM1 from the address 0, and that the program

counter is equal to 0 at the beginning of the read process. The columns a0-a3 give the values of the addresses computed from the current value of the program counter PC. The columns c(0)-c(3) give the 5 values of the codes read in the memory in the order in which they are. The columns c'(0) c'(3) give the values of the codes after reorganization by the circuit PCC. The last four columns c'(0)-c'(3) give the values of the codes after filtering in the circuit CSC.

10 If we observe the last four columns of Table 3, it can be seen that the codes therein are equivalent in arrangement and in value to the codes that were in the memory of the standard processor described in the introduction (see Table 1). Thus, the 15 reading system according to the invention reconstitutes the NOP codes which, in the prior art, were recorded in the program memory.

In general, the making of the circuits IAC and CSC is within the scope of those skilled in the 20 art, especially through techniques for the automatic designing of logic circuits using VHDL language. The algorithms A1 and A3 are expressed in VHDL language and the circuits IAC and CSC are automatically generated by a VHDL language compiler in the form of wired logic 25 circuits. A non-restrictive description shall nevertheless be given, with reference to Figures 3 and 4, of exemplary embodiments of the circuits IAC and CSC in the case where the parameters I and N are equal to 4.

30 In Figure 3, the circuit IAC receives the value of a program counter PC at input. This value is encoded on 32 bits. The result R is obtained simply by separating the two LSB wires of the physical line conveying the value of the program counter PC, with R 35 being equal to the two LSBs of the program counter. Similarly, P0 is equal to the 30 MSBs of the program counter and is obtained by separating the 30 MSB wires.

The parameter P0+1 is computed by an adder receiving the values P0 and 1 at input. The addresses a0-a3 are delivered by multiplexers MXA1-MXA3 receiving the parameters P0 and P0+1 at input. These multiplexers are
5 driven by comparators CMP0-CMP3 receiving the index i at a first input and the parameter R at a second input.

In Figure 4, the circuit CSC is made by four multiplexers MXB0-MXB3 each receiving a code $c'(j)$ at a first input and an NOP code at a second input. Each
10 multiplexer is driven by a validation signal val0, val1, val2, val3 whose value determines the selection of the code $c'(j)$ or the NOP code at output of the multiplexer. The validation codes are generated by a cascade of AND gates forming a daisy chain type of
15 chain. Each AND gate receives, at input the validation signal for the previous code and the parallelism bit of the code received by the multiplexer to which it is related. Thus, the presence of a parallelism bit at 0 in a code $c'(j)$ causes all the following validation
20 signals to change over so that the concerned code and all the following codes are replaced by NOP codes.

The validation signal val(0) of the first code $c'(0)$ is generated by an inverter gate because the parallelism bit of the first code, in principle, is
25 always equal to 0 (/p). The multiplexer MUXB0 and the selection signal val0 of the first code $c'(0)$ are, in principle, unnecessary since the first code $c'(0)$ is, in principle, never filtered, and is a large instruction comprising at least one instruction code.
30 However, the designing of this multiplexer is used to detect a possible defect in the reading of the program memory or a malfunction of the shift circuit PCC which appears when the parallelism bit of the first code $c'(0)$ is different from 0. In this case, the validation
35 signal val0 may be used as an error signal ERR.

The circuit CSC also comprises an adder ADD2 receiving the four validation signals val0 to val3 at

input. The output of this adder delivers the value n which is applied to the adder ADD1 for the incrementation of the program counter (Figure 2). The increment n is thus equal to the number of enabling 5 signals equal to 1 and consequently represents the number of instruction codes that have not been filtered, namely the effective number of codes included in the large instruction that has just been read.

The present invention is, of course, open to 10 several variations and embodiments. In particular, it is possible, in practice, to consider incorporating the filtering circuit CSC into the instruction decoder IDEC. Indeed, the function of this circuit is to ensure that codes not belonging to an instruction to be 15 executed will not be sent to the execution units EU0-EU3. The filtering of the undesirable codes may thus be done downline from the instruction register IDEC by the decoder IDEC itself. In general, the positioning of the circuit CSC in the pipeline stages is dictated by the 20 timing constraints relating to access to the memory PMEM1 and by the decoding times in the decoder IDEC, and come within the scope of those skilled in the art.

It will be clearly seen by those skilled in 25 the art that the present invention provides a wide degree of liberty in the designing of a VLIW type processor. It is especially possible, as indicated above, to design a processor having a number I of memory banks and a number N of execution units where the number I is different from the number N. It is also 30 possible to make a dynamic choice of the number N of active execution units, depending on the application planned, by inhibiting certain execution units to limit the current consumption. This inhibition of certain execution units may be designed to be programmable. The 35 architecture of a processor according to the invention of this kind is not permanently fixed, like that of a standard VLIW processor. This is due to the

elimination of the obligation to read N instruction codes en bloc and the obligation resulting therefrom to inject NOP codes into a program when an instruction contains fewer than N codes. The present invention
5 offers a wide range in the designing of a new type of processor reconciling the requirements of low consumption and compactness in terms of silicon surface area, and providing high performance characteristics through a high degree of parallelism.

APPENDIX

Reminders:

$$I = 4; N = 4$$

$P0 = \text{Quotient PC/I}$

$R = \text{PC}[\text{modulo } I]$

$R' = R$ of the previous clock cycle

$a_i = P0+1$ if $i < R$, else $a_i = P0$

$c'(j) = c(i)$, with $i=(j+R')[\text{modulo } I]$

TABLE 3: Reading of instructions INST1 to INST3 in memory banks M0 to M3

Cycle	PC	R	R'	P0	a0	a1	a2	a3	c(0)	c(1)	c(2)	c(3)	(Outputs of banks M1-M3)				(PCC circuit outputs)				(CSC circuit outputs)			
													(Addressing of banks M1-M3)				(PCC circuit outputs)				(CSC circuit outputs)			
1	0	0	-	0	0	0	0	0	-	-	-	-	c0	c1	c2	c3	c0	c1	c2	c3	c0	NOP	NOP	NOP
2	1	1	0	0	1	0	0	0	c0	c1	c2	c3	c0	c1	c2	c3	c0	c1	c2	c3	c0	NOP	NOP	NOP
3	3	3	1	0	1	1	1	0	c4	c1	c2	c3	c1	c2	c3	c4	c1	c2	c3	c4	c1	c2	c3	c4
4	7	3	3	1	2	2	2	1	c4	c5	c6	c3	c3	c4	c5	c6	c3	c4	c5	c6	c5	c6	c5	c6
5	8	0	3	2					c8	c9	...	c7	c7	c8	c9	...	c7	c8	c9	...	c9	NOP		